Instrumentation amplifiers and thermocouples

In this lab we will build and explore the basic instrumentation amplifier which is used for measuring very small voltage differences. As an example, we will measure temperature using a thermocouple. We we supposed to discuss this circuit and thermocouples in lecture, but it snowed instead. For now, you can also consult the source of all knowledge for some further details.

http://en.wikipedia.org/wiki/Instrumentation_amplifier

http://en.wikipedia.org/wiki/Thermocouple

We will discuss this further in lecture on Wednesday.

Part 1: Building, testing, and characterizing the amplifier

The circuit that you will build this week is shown below in Figure 1. We will discuss this circuit in class, and again in lecture on next Wednesday.

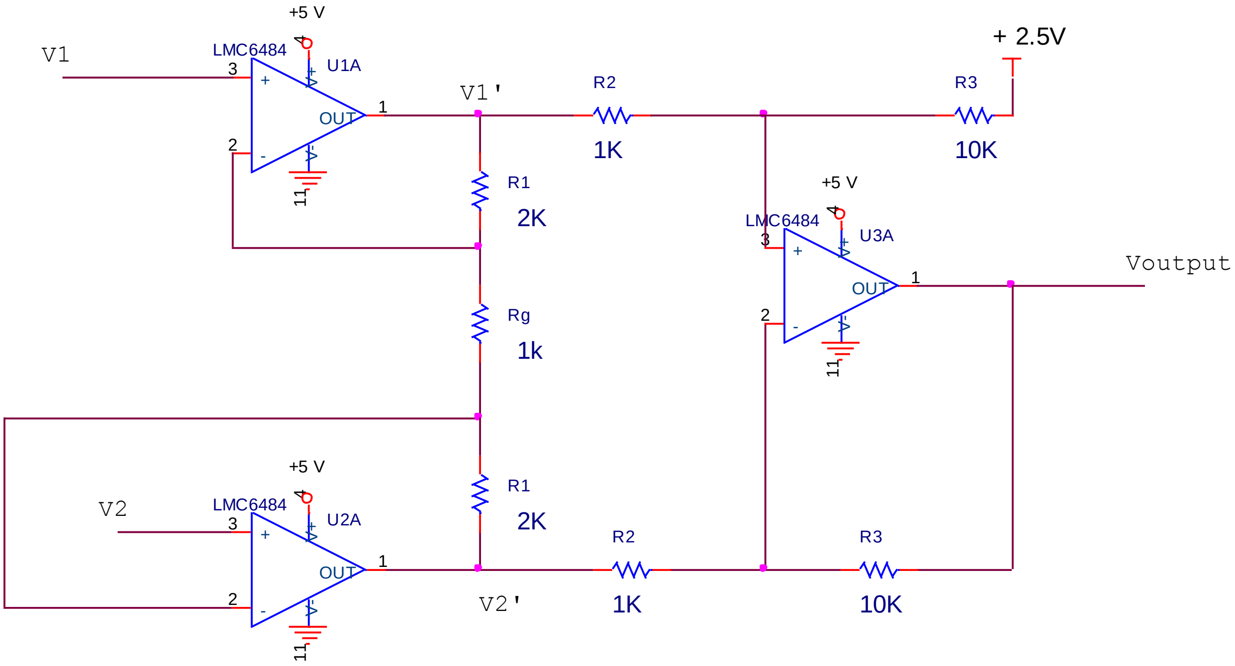

Figure 1: Instrumentation amplifier. Note that in the schematic, the pins for the input and the output may vary as you will be using a quad chip – four op-amps on one chip.

Instead of building the circuit all at once, we will build it in stages testing different parts as we go. This will help us understand what each part is doing and will allow us to test and debug the circuit. This circuit will not work well or at all if it is built poorly. The amplifier will have a relatively large gain and will be susceptible to noise. Long loopy wires will cause very poor results. Also, given that the circuit has three op-amps and several resistors, a little patience will pay off in the end. Build slow and carefully. Keep the wires short and flat to the breadboard. Clip the resistors. Make it beautiful, it will work better.

Start by building the first stage of the amplifier, shown in Figure 2, which act as buffers for the inputs, V1 and V2.

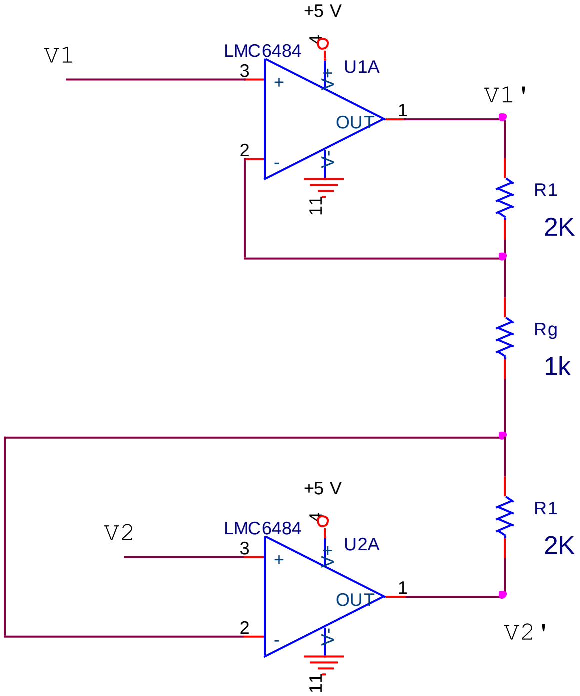

Figure 2: First stage of instrumentation amplifier.

You should be able to analyze this circuit easily and quickly determine the expected voltages at V1' and V2' in terms of V1 and V2. Do this analysis and write down the answer you expect.

Use your DAQ and a simple program to test this behavior. You can start with the template file on the website for testing this circuit. This template will output a voltage from the analog output channel on your DAQ, AO0. The analog output should be wired into V2 on the circuit in Figure 2. V1 should be connected to the 2.5 V reference on your protoboard. You can then measure the voltages V1' and V2' with respect to ground using AI0 and AI1 in differential mode. Wire V1' and V2' into the AI0+ and AI1+ respectively. Wire AI0- and AI1- into ground on your protoboard. Run the program and test that the voltages you get at V1' and V2' agree with your expectation.

Note that the voltages you are using as inputs are centered around 2.5 V. We center around 2.5 because the op-amp is supplied with 0 and 5 Volts. If we wired the above circuit such that V1 was connected to ground and V2 was a value above 0, then the circuit above would not work since the op-amp cannot output anything less than 0 Volts.

For your lab report you need to include your simple analysis and a table which includes your measured vs. expected values for V1' and V2', the difference between V1' and V2', and the gain (V2'-V1')/(V2-V1) of this component of the circuit. You can sample the voltages for a second and use the mean value. Vary the value of the voltage difference you are inputting from 20 mV up to 500 mV. You can either vary the voltage manually and rerun the program a few times, or you could adjust the script to automate the process in a loop. You can report your results with a graph or a table, whatever you like. You should notice the experiment matches your expectation better at higher voltages.

Once you believe the circuit is working, continue to add the rest of the amplifier as shown in Figure 1. Note that in Figure 1, the difference amplifier is referenced to 2.5 Volts; this feature is so that the output will fall in the middle of the voltage supplied to the op-amp.

After you build the entire circuit, leave the inputs V1 and V2 as they were previously (set to AO0 and the 2.5 V reference). Connect the output of the amplifier to AI0+ and connect AI0- to the 2.5 V reference on your protoboard. Now, rerun your program and measure the output. If everything is working, you should find that the amplifier is giving, Vout = 50 (V2-V1). As we discuss below, if the gain is off by 10% or so, everything is probably correct. You should analyze the circuit to convince yourself that the gain should be 50.

Experimentally, test the overall gain at a few conditions by varying the voltage your program puts out at AO0. Note that the op-amps will saturate when their output exceeds the 0 and 5 volt rails. The largest difference your circuit can put out is about 2.5 V – the highest voltage difference you can input is about 50 mV. While you should get a typical value close to 50, it is not uncommon for this value to be off by a fair amount.

For your lab report, you will need to analyze the error in setting the gain for your amplifier due to imprecise resistors. We will discuss the analysis more in the Wednesday class. To analyze the circuit, just use the golden rules (no current goes into the op-amp's inputs, and the op-amp inputs are equal when wired with negative feedback). To do this analysis, analyze the circuit in Figure 1, but assume that the paired resistors labeled R1, R2, and R3 are not exactly equal. Re-label the paired resistors R2A and R2B (for example). For simplicity of this step, you can assume that the resistor R3 on the top of the circuit is connected to ground (not 2.5 V). You should obtain an expression which has an extra term than previously appeared. If your analysis is correct, this term should vanish when the resistors R1, R2 and R3 are exactly matched. The reality is that the resistors are not perfectly matched. They have a stated specification of being accurate to 1%. To calculate the uncertainty in setting the gain, do the following:

1) Write a short matlab program that calculates the predicted gain for general values of R1A, R1B, R2A, R2B, R3A, R3B, and Rg. This is a prediction, not an experiment!
2) Run the program 7 times, each time increasing only 1 resistor's value by 1%. Each time you run, set all but 1 resistor's value to the nominal values in Figure 1. Each time record the value of the calculated gain for the imperfect circuit minus the nominal gain of 50. Call this difference the error; E1, E2, E3, E4, E5, E6, and E7.

3) Take all 7 errors, square each of them, sum them, and then take the square root. $\sqrt{E1^2 + E2^2 + E3^2 + E4^2 + E5^2 + E6^2 + E7^\wedge 2}$. The value is the total uncertainty in the gain for any given circuit. Compare your measured gain to the nominal gain of 50 in light of the possible uncertainty due to imperfect resistors.

Note that the error analysis above is not fully complete. A complete analysis of the uncertainty would have to incorporate the specifications of the op-amp and DAQ. However, the above analysis is correct for this one aspect of the system.

If this analysis is unclear, you should seek help from the course staff.

Part 2: Thermocouples

Now we will replace the input to our instrumentation amplifier with a thermocouple to measure temperature. Thermocouples use wires of dissimilar metals to generate a small voltage when held at different temperatures. We will use what is known as a type J thermocouple which has an iron and constantan wire. We will use two thermocouples to measure a temperature difference between the locations where the junction is made.

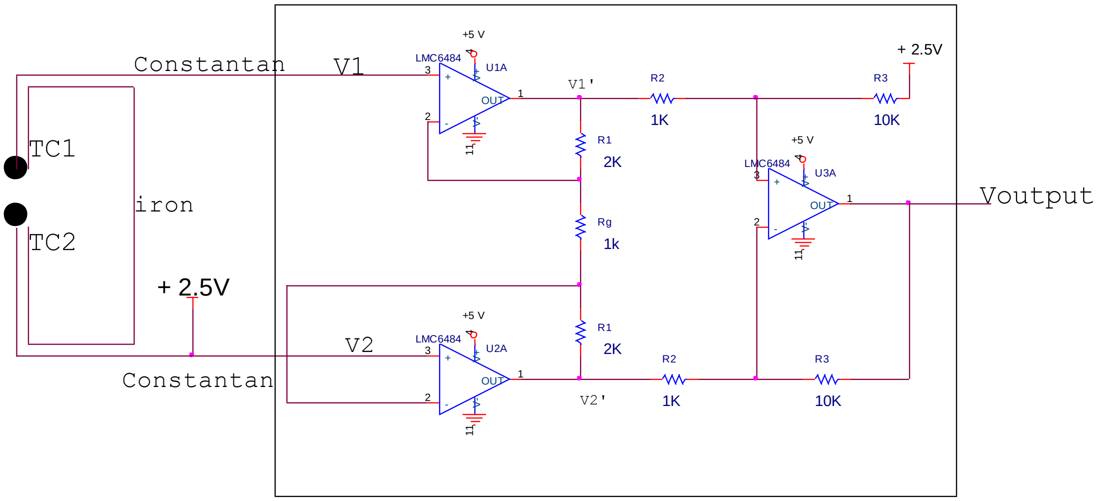

Figure 3: Thermocouple amplifier.

Treat the circuit you have already built as a component which you have already tested. You can hook the thermocouple directly to this amplifier circuit. Take two lengths of type J (iron-constantan) thermocouple wire. Strip some wire at both ends. At one end of each wire, twist the iron and constantan tightly together in order to make two thermocouple junctions. Wrap the wire so the electrical connection is very good, use pliers to twist and clamp. At the other end, twist the iron wires of the two thermocouples together. Now the remaining two constantan wires go into your instrumentation amplifier. Since the thermocouple voltage is floating, anchor it to the 2.5 V reference as shown.

Now the voltage that you measure across the thermocouple junctions will be related to temperature difference. Data for the type J thermocouple is found here:

http://srdata.nist.gov/its90/download/type_j.tab

This calibration, done by NIST, is conducted where one of the thermocouples is held fixed at 0 C. Fixing a point at 0 C is easy to do to very high accuracy. A slurry of water and crushed ice will very reliably be 0 C. Fortunately, you have an endless supply of such substance.

Start with the template code on the course website, the thermocouple calibration is already done for the type J polynomial fit, assuming one of your junctions is at the ice point. Test the program. If both thermocouples are in free air, you should see 0 V and 0 C. However, it is likely that you will register something on the order of a fraction of mV. This error corresponds to a significant temperature error. You can likely manipulate the "temperature" a fair amount by simply moving around your circuit (but not touching). The circuit is so sensitive that you influence the amount of interference and noise by simply moving around.

If you are ambitious you can add a simple low-pass filter (1K and 10uF) followed by an op-amp buffer between the output of your instrumentation amplifier and your DAQ. This is not required but will make your signal less noisy.

Get an ice bath and plunge one of the thermocouples into the ice bath (you want a slushy ice and water mixture). Make sure the thermocouple is submerged in the ice bath and that it is not near the walls of the cup. One junction will be assumed to be the ice reference based on whether it is the positive or negative input to the amplifier, so when you plunge one thermocouple junction in ice, you should see the temperature computed by the program going up. In a perfect world, this temperature will go up to 20 C or so. It is likely you will be in an imperfect world. Record the temperature of the room and comment on how close you get to room temperature in your lab report. **Please do not leave cups of water/ice lying around. Dispose of them when you are done.**

Make one measurement of the temperature of a cup of hot water (relative to the ice bath) and compare your thermocouple measurement compared to the laboratory standard (we will have one in AC428). This hot water measurement might be more accurate than room temperature measurements since the voltage is a little higher. Comment on how close your result is to the lab standard.

Finally, make a temperature measurement of one thing that is interesting to you. This can either be a single measurement or the temperature change in time. Note that the polynomial calibration in the template program we supply is only good for 0-760 C. The type J thermocouples go from -210 to 1200 C. You can use the NIST tables or calibration (see link above) if you go outside the 0 to 760 C range. Things that might be interesting are: How cold is the bottom of a snow bank? How hot is a match? How hot is my hair dryer? How hot is my car exhaust and how long does it take to warm up on a cold day? How long does a cup of coffee take to cool? How hot is my computer? How hot is a soldering iron? What is the temperature difference from inside to outside a pane of glass (related to the R-value)?

Basically anything is fine – just use your common sense and don't try to measure anything dangerous. If you find yourself wondering if what you are doing is a stupid thing to do, it probably is. **If you need to remove a DAQ from the lab, please only do so for the short time you need the test and return it promptly. We do not have enough equipment for it to go missing. Not returning equipment is theft. Do not take any equipment from the lab during the normal lab times (Tue/Fri 9-12:30)**.

Finally, in some case using the ice reference is inconvenient. If you have two thermocouples you get a measure of the temperature difference. Unless the temperature range is very large, the thermocouple is reasonably linear and thus two thermocouples can give a reliable difference, even if one point is not known. For example, if you see the NIST tables the EMF voltage between 0 and 100 C is close to the difference between 100 and 200 C. You can also do without the ice reference and use a single thermocouple if you know the temperature of the point you contact your thermocouple with the circuit. This is called cold junction compensation. You can do this "by hand" to your data after the fact in an approximate way if you need to. All you need to know is the approximate temperature of the circuit board. Further details can be found here: http://www.omega.com/temperature/Z/pdf/z021-032.pdf